DRAFT

Heavy Ion SEE Test Report for the Samsung 4Gbit NAND Flash Memory for MMS

Tim Oldham, Perot Systems Government Services, Inc. /NASA-GSFC
Mark Friendlich, MEI Technology Inc./NASA-GSFC
Anthony B. Sanders, NASA GSFC
Hak Kim, MEI Technology Inc./NASA-GSFC
Melanie Berg, MEI Technology Inc./NASA-GSFC

Test Date: 23 May 2009 Report Date: 18 June 2009

I. Introduction

This study was undertaken to determine and compare the susceptibilities of 4 Gbit NAND Flash memories from Samsung to destructive and nondestructive single-event effects (SEE) for the NASA MMS mission. The devices were monitored for SEUs, errors from individual cells, for SEFIs, errors arising in the control logic, and for destructive events, including latchup, induced by exposing them to a heavy ion beam at the Texas A&M University Cyclotron. Previous testing had revealed a possible sensitivity at high temperatures, leading to destructive failures of the write circuit. For this reason, high temperature testing was a major focus in this test, along with filling in some gaps in the results from the previous test.

II. Devices Tested

We tested a combine total of 23 Samsung parts in both test runs, from five different date codes (part number K9F4G08U0A-PCB0, Lot Date Codes (LDC) 840, 843, 846, 901, and 907). We began with eight parts from each date code delidded and fully operational on the bench at Goddard, that were shipped to TAMU. However, not all of them worked properly at TAMU. What we actually used was three parts from LDC 840 in the first test, and two more in this test, five parts from LDC 843 in the first test, and two more in this test, two parts from LDC 846, plus two more in this test, and four parts from LDC 901, plus two more in this test, and one part from LDC 907. These parts are listed in Table II, which is discussed in the results section below. The parts have 512Mx8 organization with large blocks. That is, the blocks are 128Kx8, with 64 pages/block. Each page is nominally 2Kx8, but they also have 64 redundant columns, which makes the total page size 2112x8. NAND flash normally has some bad blocks which can be screened off. The specification is that no more than 80 of the 4096 blocks will be bad. In our experience, the parts almost always have a few bad blocks, but it is usually a single digit number. Note that with commercial devices, the same lot date code is no guarantee that the devices are from the same wafer diffusion lot or even from the same fabrication facility.

The device technology is 63 nm minimum feature size CMOS NAND Flash memory. All the parts are single die, SLC (single level cells). The chips came in a 48-pin TSOP package, but the plastic had been dissolved on the topside to expose the chips, allowing the beam to reach the chip surface.

Fig. 1. Photos of die

III. Test Facilities

Facility: Texas A&M University Cyclotron
Flux: (5 x 10^3 to 1. x 10^5 particles/cm^2/s).
Fluence: All tests were run to 1E3 to 1E8 p/cm^2, or until destructive or functional events occurred.

Table I: Ions/Energies and LET for this test

TAMU Ions	Energy/ AMU	Energy (MeV)	Approx. LET incident on die surface (MeV•cm^2/mg)	Angles	Effective LET
Ne	15	300	2.7	0, 45	2.8, 3.9
Ar	15	600	8.4	0,45	8.4, 11.8
Kr	15	1260	30.1	0, 45	29.3, 41
Xe	15	1965	54.8	0, 45	53.9, 75
Au	15	2955	87.5	0	87.5

IV. Test Conditions

Test Temperature: Room Temperature for SEU, 70° C for SEL, 40-70º C for other high temperature destructive events.
Operating Frequency: (0-40 MHz).
Power Supply Voltage: (3.3V for SEU and SEFI, 3.6V (3.3+10%) for SEL). Standard test methods for SEU testing (e.g., ASTM 1192) call for testing at nominal voltage less 10%, because SEU in standard volatile memories is caused by voltages being pulled down. However, flash memories are designed to retain information even at zero volts, so the upset mechanisms are clearly different, here. In addition, we are also looking for control logic errors, which are thought to be due to things turning on when they are not supposed to be on. Reduced voltage would cause an underestimate of the rate for these events. Therefore, we used nominal voltage, 3.3 V, in all tests.

V. Test Methods

Because Flash technology uses different voltages and circuitry depending on the operation being performed, testing was performed for a variety of test patterns and bias and operating conditions.

Test patterns included all 0's, all 1's, checkerboard and inverse checkerboard. In general, all zeroes is the worst-case condition for single bit errors. For a zero, the floating gate is fully

charged with electrons. An ion can have the effect of introducing positive charge, which may be enough to cause a zero-to-one error. However, a checkerboard pattern (AA) was used in most of the testing because errors in the control circuitry can cause errors of both polarities. One-to-zero errors are an indication that the errors are coming from the control circuits. Between exposures, all patterns can used to exercise the DUT, to verify that it was still fully functional. However, all patterns are not used on every shot, just because it is time consuming to do so. The maximum clock frequency for these devices was 40 MHz, which is also the frequency used in the dynamic testing.

Bias and operating conditions included:

1) Static/Unbiased irradiation, in which a pattern was written and verified, and then bias was removed from the part and the part was irradiated. Once the irradiation reached the desired fluence, it was stopped, bias was restored, and the memory contents were read and errors tallied.

2) Static irradiation, which was similar to unbiased irradiation, except that bias was maintained throughout irradiation of the part.

Note that these conditions provide no opportunity to monitor functional or hard failures that may occur during the irradiation. It was also not possible to monitor the power supply current during the unbiased tests, but this was done in all the other test modes.

3) Dynamic Read, in which a pattern was written to memory and verified, then subsequently read continuously during irradiation. This condition allows determination of functional, configuration and hard errors, as well as bit errors. In this mode, the number of static bit errors is determined by reading the memory again, after the beam is turned off.

4) Dynamic Read/Write, which was similar to the Dynamic Read, except that a write operation is performed on each word found to be in error during the previous Read.

5) Dynamic Read/Erase/Write, which again was similar to the Dynamic Read and Read/Write, except that a word in error was first erased and then rewritten. In this mode, the words that are read are compared to an "expected" pattern, which is actually the complement of the stored pattern. For this reason, every word is erased, as if it were in error. Because the Erase and Write operations use the charge pump, it is expected that the Flash could be more vulnerable to destructive conditions during these operations.

6) In the previous test, latchup (SEL) testing was conducted at 70° C, and 3.6 V. There were no cases where SEL was observed, but there were other destructive failures at high temperature. In this test, we did extensive testing in the range 40-70º C, although we did not consider it to be SEL testing. The goal of this high temperature testing was to detect destructive events, other than SEL.

7) In this set of experiments, we have attempted to look at angular effects, which may include multiple bits grazed by the same ion, and other effects due to charge sharing by multiple nodes in the control logic. This test was done with at 45 degrees, which was close to the maximum possible angle, because the socket would have blocked the beam at angles much higher. There were two orientations, which we referred to as 45º North, and 45º East. Although the normal bit error upset rate is somewhat higher at high angles, probably because of charge sharing in the control logic, destructive

failures occur primarily at normal incidence. For this reason, much of the high temperature testing was done at normal incidence.

The Block diagram for control of the DUT is shown in Figure 2. The FPGA based controller interfaces to the FLASH daughter card and to a laptop, allowing control of the FPGA and uploading of new FPGA configurations and instructions for control of the DUT. Power for the flash is supplied by means of a computer-controlled power supply. The National Instruments Labview interface monitors the power supply for over-current conditions and shuts down power to the DUT if such conditions are detected.

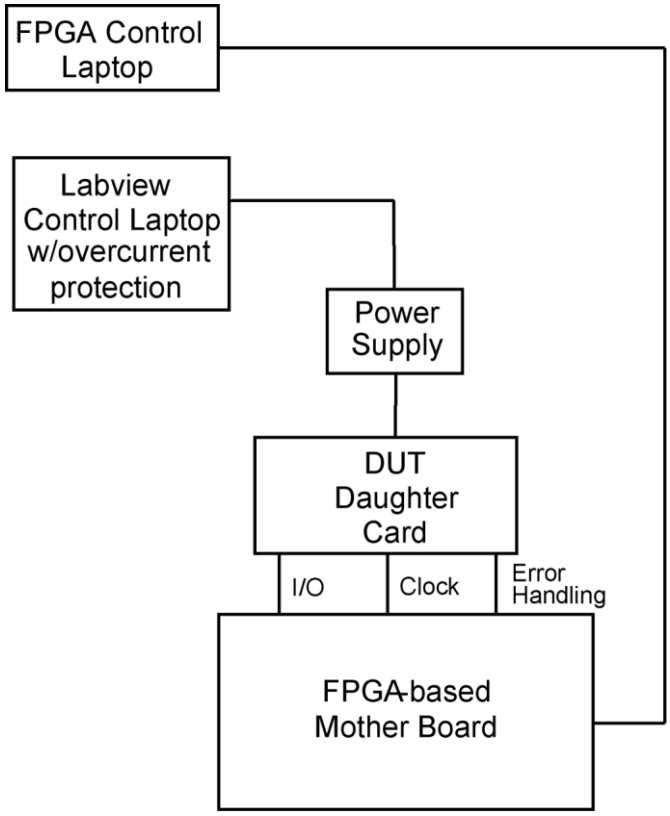

Figure 2. Overall Block Diagram for the testing of the NAND Flash.

Fig. 3. Front and back views of motherboard and daughterboard, with DUT.

VI. Results

During testing, the DUTs were irradiated with the ions indicated in Table I. The DUT was oriented normal to the incident beam, or at 45 degrees. The errors observed in static SEU testing are shown in Fig. 4, with no bias applied. The 45 degree data is plotted at the effective LET (LET/cos θ). This is done so that one can distinguish between the normal incidence shots and

the 45 degree shots. It is not done because effective LET is expected to be a useful concept for other reasons.

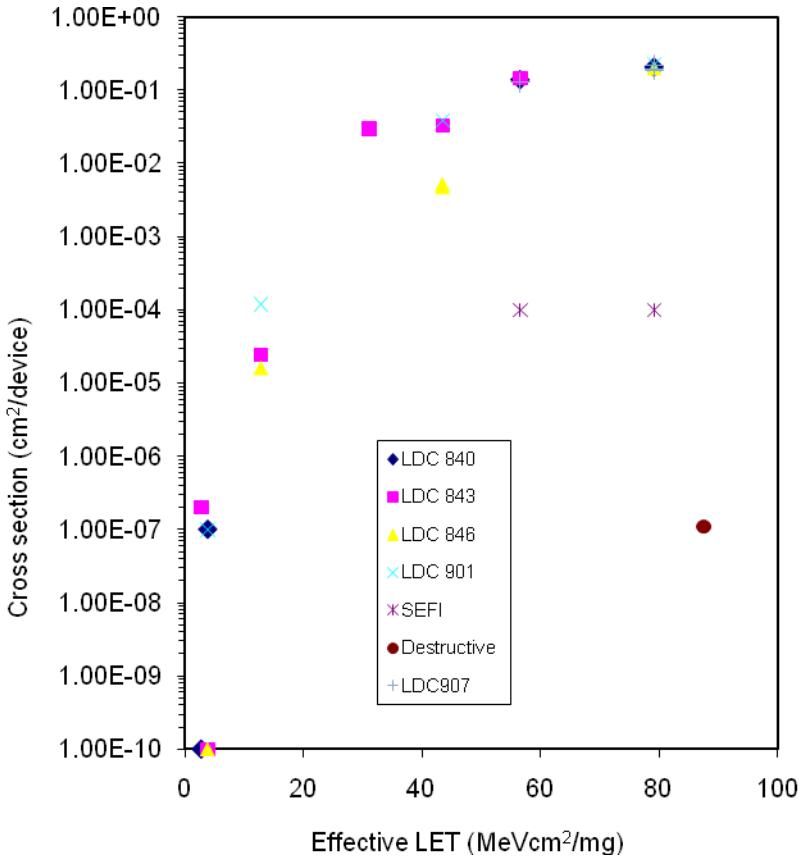

Fig. 4. SEU, SEFI, and Destructive results in static, unbiased test mode.

Fig. 4 shows new results in the unbiased static mode for LDC 907, and also for LDC 840 with Xe ions, both for normal and high angle incidence. LDC 907 had not been tested at all previously, and LDC 840 had been tested only at very high and very low LET, with Au ions and Ne ions (LET = 2.8, and 87). For Xe, LET=56 at normal incidence, and LET_{eff}=79 at 45º, which adds two intermediate points for LDC 840. LDC 907 was only tested with Xe. The results in Fig. 4 for these two LDCs are virtually identical, and there is very little difference between any of the LDCs in this test mode, at any LET. The one destructive event occurred with Au ion, (LET = 87).The results using the two different orientations, 45º E and 45º N, were that the total cross sections were essentially the same. For this reason, we did most of the high angle exposures in the 45º E orientation, which was easier to align. However, when we did the post-processing of the output files after the run was over, and examined the bit error maps in detail, there was a difference in the error signature. When the beam came from the "North", there were often pairs of errors, where the same column address would have the same bit in error in two successive pages within a block, for example. Or successive pages would have errors differing by two columns, with the same bit in error at both addresses. Apparently, having the beam incident along the columns causes more of these errors than when the beam angle is across the columns.

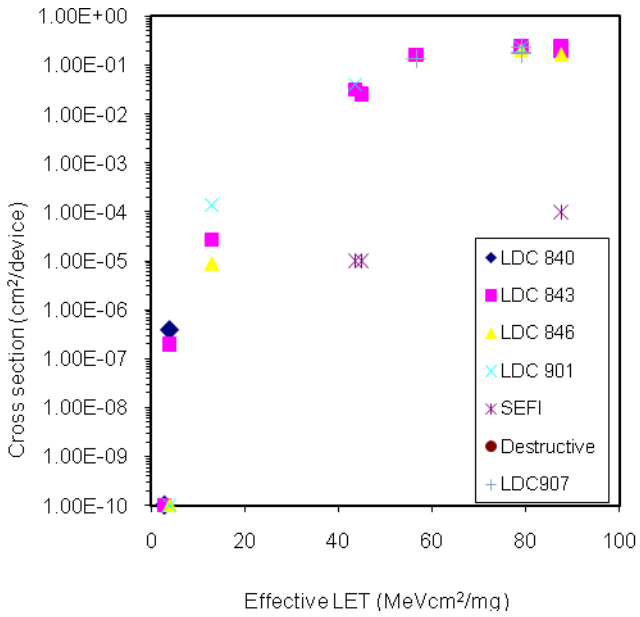

Fig. 5. SEU, SEFI, and Destructive results for static mode, with bias applied.

The results for static mode with bias applied are shown in Fig. 5. The main effect is single bit upsets, along with a few SEFIs, but no destructive events. Again, the results for LDC 840 and LDC 907 are very consistent, with each other, and with the other LDCs. They are also very consistent with the unbiased results in Fig 4, except for the one destructive event with Au ions. In both Figs. 4 and 5, only one shot was at high temperature, 70º C, with Xe ions at 45º, in both cases. There were no SEFIs, no destructive events, and the bit error rate was consistent with shots under other test conditions.

For the Dynamic Read condition, the parts showed exhibited transient read errors in addition to the bit errors, which are plotted in Fig. 6. In this mode, the DUT reads continuously with the beam on. The significance of the transient errors is not always completely clear, because the entire memory can be read multiple times, which means static errors will be read multiple times. In addition there are errors due to transient noise in the read circuit or the control logic. The static errors are bit errors read after the beam is turned off. In this test mode, there were two SEFIs and no destructive events. The bit error cross-section is similar for all five LDCs. Again, there was only one high temperature shot in this test mode, for LDC 907, with Xe at 45 degrees. There was no SEFI, and the bit error cross-section was consistent with the other shots.

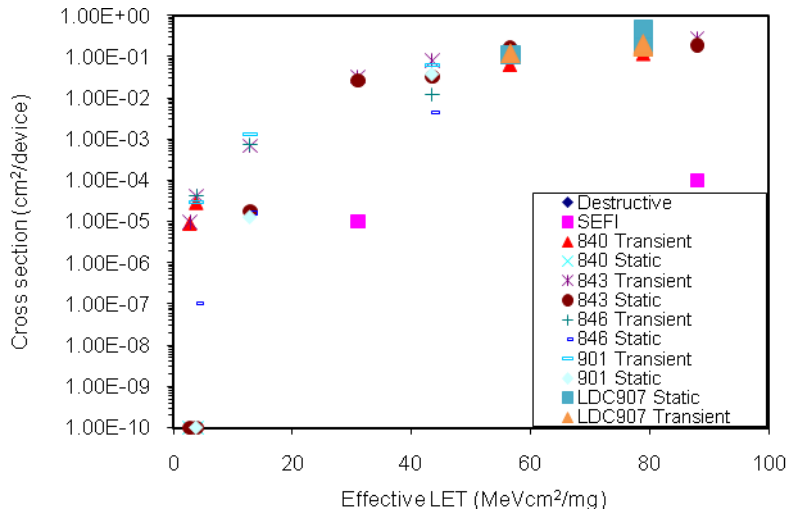

Fig. 6. SEU,SEFI, and Destructive results for Dynamic Read mode.

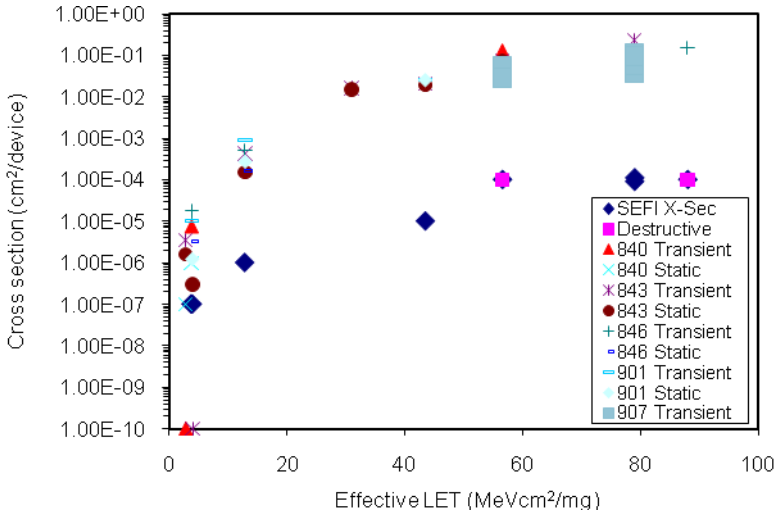

Fig. 7. SEU, SEFI, and Destructive results in the dynamic R/W mode.

Results of the dynamic R/W tests are shown in Fig. 7. Generally these results are unremarkable, because the usual zero-to-one errors are rewritten as they occur. For this reason, there are fewer errors indicated than in Fig. 6, although the difference is not large. The main reason for including this test was the expectation that the high voltage write operation would contribute to more errors in the control circuits, which appears to have happened—there are seven SEFI events, and two destructive failures. One of these was an erase failure with Au, and was previously reported. The other destructive failure, a write mode failure for LDC 840, occurred at normal incidence and room temperature. Where a static cross section is given, it is based on the number of errors detected after the exposure and resetting of the DUT, as before. The transient cross section is based on errors detected during the exposure. But some of the transient errors are probably really static bit errors that were rewritten during the test. The SEFIs were cases where the DUT stopped responding to commands, until power was cycled, or there

were large numbers of non-random errors. In each case except the two destructive events, the DUT was restored to normal operation and used for the next shot. There were nine shots at elevated temperature in this test mode, all but one at 45º, and no destructive failures. These will be discussed in more detail later.

Results for the dynamic R/E/W tests are shown in Fig 8. In this mode, errors are counted as they are read, but then they are erased and rewritten. Therefore, there are no static errors read after the beam is turned off, and bits in error for a time are counted as transient errors. Because the high voltage erase and program operations are performed constantly, it is expected there will be errors in the control logic in this mode. In fact, there were three SEFIs and two destructive failures at room temperature in this mode, and many more at high temperature. In the earlier March test, one SEFI and one destructive event occurred with Ar ions incident at 45 degrees (LDC 0843). Immediately after the shot, the DUT did not respond to commands. After power was cycled, the part responded to commands, except that the write circuit had failed. With Au ions at normal incidence, there was also a SEFI with a destructive failure (LDC 901), also in the March test. This shot was counted as a SEFI because of multiple block errors, but the erase circuit also failed. In the May test, there were 55 shots in this test mode, with all but five at elevated temperature. These are included with the room temperature results in Fig 8. There were six destructive events at high temperature, which were all write failures. Room temperature failures, on the other hand, are usually erase failures. In Table II, we summarize all the destructive events, most of which were at high temperature.

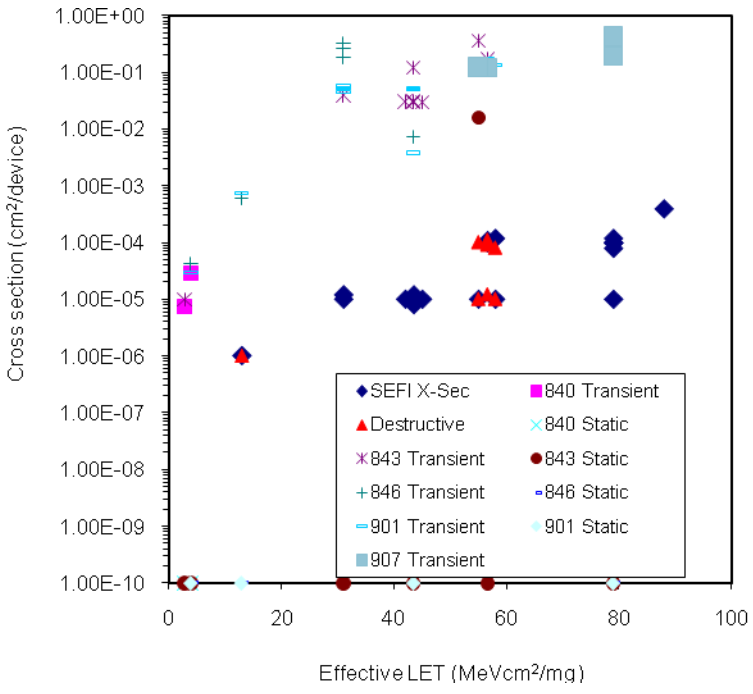

Fig. 8. SEU, SEFI, and Destructive results in Dynamic Read/Erase/Write mode.

Table II. Summary of Samples Used and Failure Modes.

LDC	Sample No.	No. Shots	Ion	Fluence (p/cm²)	Angle (degrees)	T (º C)	Comments
840-March	2	1	Au	9e6	0	25	Unbiased static, erase failure
	3	10	Ne	1e7	0	25	1 SEFI, still functional
	3	1	Xe	1e5	0	70	R/E/W—write fail
	4	1	Xe	1e4	0	70	R/E/W—write fail
840-May	B7	10	Xe	1e4/shot	5 @0, 5@45	25	Write fail in R/W mode, normal inc.
	B6	1	Xr	1e4	0	25	Write fail, R/E/W
843-March	1	3	Au	1.2E5 total	0	25	Erase fail—Dyn. Read
	2	10	Ne	1e7/shot	5@0, 5@45	25	1 SEFI, but still functional
	7	10	Xe	1e4/shot	5@0, 5@45	25	1 SEFI, but still functional
	3	10	Kr	1e5/shot	5@0, 5@45	25	1 SEFI, but still functional
	4	5	Ar	1e6/shot	5@45	25	Erase fail—R/E/W
843-May	20	12	Xe	1e4/shot	11@45E, 1 normal	40-70 40	Normal inc. fail from stuck bits (incomplete write)—R/E/W
	22	6	Kr	1e5/shot	2@45E, 4@0	40-70 40-70	1 SEFI, but still functional
846-March	4	2	Au	1e4/shot	0	25	R/W—write fail
	7	5	Ne	1e7/shot	45	25	No SEFIs, fully functional
		5	Xe	1e4/shot	45	25	No SEFIs, fully functional
		6	Kr	1e5/shot	45	25	3SEFIs, fully functional
		5	Ar	1e6/shot	45	25	No SEFIs, fully functional
846-May	B10	4	Xe	1e4/shot	3 @ 45, 1@ 0	40-70 40	Write fail at normal inc., R/E/W
	B12	12	Kr	1e5/shot	4 @ 45, 4 @ 0	40-70 40-70	4 total SEFIs, but fully functional
901-March	1	2	Au	3.5e4 total	0	25	Write fail—R/E/W

	3	5	Ne	1e7/shot	45	25	No SEFIs, fully functional
		5	Xe	1e4/shot	45	25	No SEFIs, fully functional
		1	Xe	1e5	45	70	SEFI, functional after PC
	4	1	Xe	1e5	0	70	Write fail
	6	5	Kr	1e5/shot	45	25	No SEFIs, fully functional
		5	Ar	1e6/shot	45	25	No SEFIs, fully functional
901-May	B04	9	Xe	1e4/shot	4 @ 45, 5@ 0	40-70 40-70	Write fail at normal inc., 70 C R/E/W
	B03	7	Kr	1e5/shot	3 @ 45, 4 @ 0	40-70 40-70	2 SEFIs, fully functional
907-May	13	6	Xe	1e4/shot	0	25	No SEFI, fully functional
		10	Xe	1e4/shot	45	25	No SEFIs, fully functional
		6	Xe	2.4e5 total	45	70	1 SEFI, fully functional
		6	Xe	1e4/shot	45	2 ea. @ 40, 50, 60	1 SEFI, fully functional
		3	Xe	1e4/shot	0	2 @ 40, 1 @50	Write fail @ 50, R/E/W Write current high before normal inc. shots

In addition to the room temperature results shown in Fig. 8, there were four shots at elevated temperature (70º C) and increased voltage (3.6 V) to test for SEL, in the March test. All were using the R/E/W test mode, with Xe ions incident. In one shot, the ions were at 45º incidence, and the DUT survived at a fluence of 10^5 particles/cm^2. A SEFI was recorded, but after power cycling, the DUT was successfully reprogrammed. In this case, the DUT was from LDC 901. In three other shots at normal incidence, the DUTs all suffered write circuit failures at lower fluences (LDC 0901 and two parts from LDC 0840). For every shot, the power supply current was monitored continuously, and it is clear from the current traces when these failures occurred. The current traces typically show a read current of about 10-11 mA, with pulses up to 15-16 mA when errors are being rewritten. When the write circuit starts to degrade, the write current will increase, but parts with current in the range of 22-23 mA, or even higher, have still functioned properly. On these shots, the current suddenly increased from 22-23 mA to 30 mA, then 40, then 50 mA, or more, which meant the part had failed. This high current was not the result of SEL, because current returned to normal when the DUT was told to stop writing—no power cycle was

necessary. The DUTs all worked properly in read mode, and they could be erased, but not rewritten. Because normal incidence was clearly the worst case, our approach in the May test was to start at oblique angles, and collect as much data as possible. Since we already had 25º C data for most test conditions, we started the normal incidence shots at 40º C, and worked up. If the parts did not survive to 70º C, at least we knew what temperature they did survive to. Let us summarize the results of these tests: all the high T failures were with Xe ions (LET=56), none were observed with Kr (next lowest LET=31). Because there were no failure with Kr, no testing was done with other, lower LET ions. All the high temperature failures were at normal incidence—there was not a single failure at high angles. All the failures occurred in either the R/E/W test mode or the R/W mode—there were no failures if the high voltage erase and write operations were not being performed. Summarizing the results in Table II by LDC:

1. LDC 840 had four parts that failed in either the R/E/W or R/W test modes. Two failed on the first 70º C shot at normal incidence, in the March test. Two more failed at room temperature and normal incidence in May.

2. LDC 843, there were two erase failures at room temperature and normal incidence in March. One was with Au ions in the Dynamic Read mode, and the other was in R/E/W mode with Ar ions. In the recent May test, there was one additional failure, with Xe ions at high temperature. This sample failed on the first high temperature shot at normal incidence, at 40º C, but it had survived 11 other shots at higher angles, at temperatures over the range 40-70º C. This part had a write mode failure, but it was unlike the other write mode failures. In most cases, when the write circuit failed, it failed completely, and every address in the entire memory was bad. For this part, however, the write circuit worked for all but 159 addresses (stuck bits), even though it was drawing 53 mA, compared to 15-16 mA, nominal write current. This was an incomplete write, not a complete failure.

3. LDC 846 had one write failure at normal incidence and room temperature with Au ions in March. One other sample failed in the recent May test at normal incidence and 40º C, with Xe ions. This part had survived three previous shots at 45º incidence, and 40-70º C.

4. LDC 901 had a write failure with Au ions at normal incidence and room temperature, and another write failure with Xe ions (normal incidence, 70º C) in the initial March test. In the later May test, there was one additional failure, at normal incidence and 70º C. However, the sample had survived four previous shots at 45º incidence, 40-70º C, and also four shots at normal incidence, over the same temperature range. This sample had survived one shot at 70º C and normal incidence, with fluence of 10^4 particles/cm^2, but it failed on a second shot under these conditions, when the fluence was increased to 10^5 particles/cm^2.

5. LDC 907 had only one sample tested, because of what we believe was corrosion on the leads of several others, which made them unusable. This one sample was exposed for no less than 31 shots, all with Xe, and 15 of them at elevated temperature. It survived twelve shots at 45º incidence, 40-70º C. It also survived two normal incidence shots at 40º C, before failing at 50º C. The write current had risen on this part, from the normal 15-16 mA, to above 25 mA, before the normal incidence shots began. For this reason, we expected it to fail, before it actually did, as a result of cumulative damage. Damage on the shot where failure actually occurred was only a small part of the story, for this part.

The flux at and above the LET of Xe in geosynchronous orbit is about one particle/cm^2 per 125 years. In the March test, we had three destructive write mode failures at normal incidence, and

70º C, with Xe ions. We estimated it took about 52,000 particles/cm^2, to produce these failures. In the follow-up May test, only one sample (LDC 901, sample B04) was tested at 70º C and normal incidence. All the others failed at lower temperature, before we could get to 70º C. This one sample survived one shot with fluence 10^4 particles/cm^2, and failed part way through another exposure to 10^5 particles/cm^2. If we estimate the total fluence to failure at $3x10^4$ particles/cm^2, then we have four total failures with a mean fluence between failures a little over $2x10^4$ particles/cm^2. In geosynchronous orbit, the flux at the LET of Xe ions is about one particle/cm^2 per 125 years, which means there would be one chip failure per 2.5 million chip years. However, this flux is integrated over 4π steradians, and the angular test results clearly show that failures happen only when the incident particle is aligned just right. This failure interval needs to be multiplied by 4π, to account for the angular dependence. These failures also occur only in the high voltage operations, Program and Erase, which are estimated to have about a 2% duty cycle—certainly less than 5%. If these correction factors are applied, the interval between failures is estimated to be more than 600 million chip-years, in geosynchronous orbit. Other orbits would generally be even longer. For a system with 2000 chips, the system failure rate would be one failure per $3x10^5$ years. If we include the three parts that failed at lower temperatures, they survived about $4x10^4$ particles/cm before the failures occurred. The totals would then be seven failures in about $1.2x10^5$ particles/cm^2, or about $1.7x10^4$ particles/cm^2 between failures. That is, the estimated interval between failures is reduced by only about 15%. We note that many more shots were taken at room temperature, and also at the next lowest LET (Kr) at temperature, and similar failures were not observed, so there is a sensitivity at high temperature (and only at high temperature), with high enough LET, that had not been noted before. There were occasional destructive failures at room temperature, but these were failures of the erase circuit, and not the write circuit, with one exception. The one room temperature write failure happened when there was a watchdog error, meaning that the DUT stopped responding to all commands. After cycling power, we found that the write circuit did not work, any longer. We did not observe the current increase described above, which was characteristic of the other high temperature failures.

To estimate the error rate expected in space, given the cross sections in Figs. 4-8, we did one CRÈME96 run for geosynchronous orbit, using the following Weibull parameters: threshold LET=2.8, Width =37, exponent = 5, and saturation cross section = 7.5e-11 square microns. This curve bounds all five of the measured cross sections, with some margin in all cases. The result was a bit error rate of 6.35e-12 errors/bit-day, which is approximately five orders of magnitude better than a typical volatile memory. For a 4G memory, this is equivalent to 0.025 errors/chip-day, or about 50 bit-errors per day for a system with 2000 chips. Handling this error rate should be well within the capabilities of error-correcting software. SEFIs are more difficult to correct, but, as Figs. 4-8 show, the cross section is typically 3-4 orders of magnitude less than the bit error cross section, even on shots where SEFIs occur. However, most shots have no SEFIs, so the average cross-section is really much lower than the Figures indicate. Based only on the cross-sections in Figs. 4-8, the system SEFI rate is estimated to be .005-.05 events/day, or one event every 20-200 days. Based on all shots, the rate is perhaps an order of magnitude lower. Most of these can be corrected by cycling power, and reprogramming the corrupted portion of the memory, so the impact to the mission should be manageable. We note that the geosynchronous orbit is a more stringent environment than the planned MMS orbit, so these rates would be lower for the actual MMS orbit.

It is not clear what the underlying physical mechanism(s) are, that are causing these destructive failures. There are two models in the literature, which could be useful in explaining some of our

results. The first of these is by Brews et al., [1], who suggested that charge from the ion strike accumulates under the oxide by a process similar to funneling [2]. This accumulated charge creates a space-charge voltage, which adds to the applied voltage. The total field exceeds the breakdown threshold for the oxide, and actually blows a hole in the oxide (gate rupture). The angular dependence, that we have observed, falls out of this model very naturally. The voltage difference across the oxide is the same, at different angles. But the conducting path will be longer at oblique angles, which means the field will be lower, so gate rupture is less likely at high angles. But this model does not account for some of our results which suggest accumulated damage might play a role, and it does not account for the apparent temperature dependence. The second model [3] suggests that each incident ion creates a small damaged region. After enough ions, these regions start to overlap, and eventually a percolation path forms all the way across the oxide. This model might explain why we see signs of accumulated damage being important. But it does not explain either the angular dependence or the temperature dependence in our results. Therefore, we conclude that there is no existing model that accounts for our results.

VII. Recommendations

Our recommendation is to use these Samsung 4G parts as flash memory on MMS. All flash memory has a bit upset rate that is outstanding, compared to typical volatile memories. The reason is that volatile memories lose information when ion strikes pull down voltages, but flash is designed to retain information, even with no voltage applied. Therefore flash is typically five or more orders of magnitude better than volatile memories in upset rate, and these Samsung parts are no exception. SEFIs are a more significant problem than bit errors in advanced flash memories, and, of course, destructive events are potential show stoppers. Both of these things are much less common for these Samsung parts than in, for example, the Micron parts tested earlier. On the Micron parts, there was a SEFI on nearly every shot, and we often lost data because the DUT basically shut down. It was actually hard to determine the upset rate because there were so many SEFIs. For the Samsung parts, SEFIs also occurred, but on a much smaller fraction of the shots. For the Micron parts, destructive events happened even in static mode at room temperature, at low LET, and at oblique angles. For the Samsung parts, destructive events also happened, but they typically required high temperature, high voltage, and just the right angle of incidence. Therefore, the risk of SEFIs and destructive failures appears to be much lower in these Samsung parts, than in others, but the risk is not zero. For this reason, it will be important to have a good strategy for managing these risks. We also note that the 2009 LDCs seem to have had better resistance to destructive effects than the 2008 LDCs, so using those parts as much as possible would seem to be a good plan. Therefore, we recommend using LDC 907 and 901 as much as possible, and not using LDC 840 if it can be avoided. In terms of resistance to destructive failures, LDCs 843 and 846 fell in the middle, and were about the same as each other.

VIII. Further Test Requirements

These Samsung parts have also been tested for total dose (TID) response, and the response has been excellent, passing well past 100 krads (SiO_2). At this point the only additional testing we would recommend is failure analysis on the parts that failed in SEE testing. If we can see where the failures occurred, and what exactly failed, it might help us to understand why the failures occurred. That understanding might help us develop mitigation strategies.

IX. References

1. J.R. Brews et al., *A Conceptual Model of Gate Rupture in Power MOSFETs*, IEEE Trans. Nucl. Sci., **NS-40,** 1929 (1993).

2. F.B. McLean et al., *Charge Funneling in N- and P-Type Substrates*, IEEE Trans. Nucl. Sci., **NS-29**, 2018 (1982).

3. B.K. Choi et al., *Long-Term Reliability Degradation of Ultrathin Dielectric Films Due to Heavy Ion Irradiation*, IEEE Trans. Nucl. Sci., **NS-49**, 3045 (2002).